arXiv:1510.08393v1 [cs.LO] 28 Oct 2015

What's Decidable about Syntax-Guided Synthesis?

Benjamin Caulfield[1], Markus N. Rabe[1], Sanjit A. Seshia[1], and Stavros Tripakis[1,2]

[1]University of California, Berkeley
[2]Aalto University

Abstract. Syntax-guided synthesis (SyGuS) is a recently proposed framework for program synthesis problems. The SyGuS problem is to find an expression or program generated by a given grammar that meets a correctness specification. Correctness specifications are given as formulas in suitable logical theories, typically amongst those studied in satisfiability modulo theories (SMT) solving.
In this work, we analyze the decidability of the SyGuS problem for different classes of grammars and correctness specifications. For theories with finite domains, we give a general algorithm to solve the SyGuS problem. Finite-domain theories include the bit-vector theory without concatenation. We prove SyGuS undecidable for a very simple bit-vector theory with concatenation, both for context-free grammars and for tree grammars. Additionally, we show that for the theory of uninterpreted functions with equality, SyGuS is undecidable even for a very restricted class of grammars. Finally, we give some additional results for linear arithmetic and bit-vector arithmetic along with a discussion of the implication of these results.

1 Introduction

Program synthesis is an area concerned with the automatic generation of a program from a high-level specification of correctness. The specification may either be total, e.g., in the form of a simple but unoptimized program, or partial, e.g., in the form of a logical formula or even a collection of test cases. Regardless, one can typically come up with a suitable logic in which to formally capture the class of specifications. Traditionally, program synthesis has been viewed as a deductive process, wherein a program is derived from the constructive proof of the theorem that states that for all inputs, there exists an output, such that the desired correctness specification holds [14], with no assumptions made about the syntactic form of the program. However, over the past decade, there has been a successful trend in synthesis in which, in addition to the correctness specification, one also supplies a hypothesis about the syntactic form of the desired program. Such a hypothesis can take many forms: partial programs with "holes" [16,17], component libraries [13,10], protocol scenarios [18,1], etc. Moreover, the synthesis of verification artifacts, such as invariants [5], also makes use of "templates"

constraining their syntactic structure. The intuition is that such syntactic restrictions on the form of the program reduce the search space for the synthesis algorithms, and thus speed up the overall synthesis or verification process.

Syntax-guided synthesis (SyGuS) [2] is a recently-proposed formalism that captures this trend as a new class of problems. More precisely, a SyGuS problem comprises a logical specification φ in a suitable logical theory T that references one or more typed function symbols f that must be synthesized, along with one or more formal languages L of expressions of the same type as f, with the goal of finding expressions $e \in L$ such that when f is replaced by e in φ, the resulting formula is valid in T. The formal language L is typically given in the form of a grammar G. Since the SyGuS definition was proposed about two years ago, it has been adopted by several groups as a unifying formalism for a class of synthesis efforts, with a standardized language (Synth-LIB) and an associated annual competition. However, the theoretical study of SyGuS is still in its infancy. Specifically, to our knowledge, there are no published results about the decidability or complexity of syntax-guided synthesis for specific logics and grammars.

In this paper, we present a theoretical analysis of the syntax-guided synthesis problem. We analyze the decidability of the SyGuS problem for different classes of grammars and logics. For grammars, we consider arbitrary context-free grammars, tree grammars, regular grammars, and grammars specific to linear real arithmetic and linear integer arithmetic. For logics, we consider the major theories studied in satisfiability modulo theories (SMT) [4], including equality and uninterpreted functions (EUF), finite-precision bit-vectors (BV), and arrays – extensional or otherwise (AR), as well as theories with finite domains (FD). Our major results are as follows:

- For EUF, we show that the SyGuS problem is undecidable even for severely restricted classes of tree grammars and even regular grammars. These results extend straightforwardly for the theory of arrays. (See Section 3.)
- For BV, we show (perhaps surprisingly) that the SyGuS problem is undecidable for the classes of context-free grammars and tree grammars. (See Section 5.)

Theory \ Grammar Class	Regular	Tree	Context-free
Finite-Domain	?	D	U
Bit-Vectors	?	U	U
Arrays	U	U	U
EUF	U	U	U

Table 1. Summary of main results, organized by background theories and classes of grammars. "U" denotes an undecidable SyGuS class, "D" denotes a decidable class, and "?" indicates that the decidability is currently unknown.

– For arbitrary theories with finite domains (FD) defined in Section 2, we show that the SyGuS problem is decidable for tree grammars, but undecidable for arbitrary context-free grammars. (See Section 4.)

See Table 1 for a summary of our main results. In addition, we also consider certain restricted grammars specific to the theory of linear arithmetic over the reals and integers (LRA and LIA), as well as bit-vectors (BV) where the grammars generate arbitrary but well-formed expressions in those theories and discuss the decidability of the problem in Section 6. The paper concludes in Section 7 with a discussion of the results, their implications, and directions for future work.

2 Preliminaries

We survey basic definitions and notation on logical theories, syntax-guided synthesis, and grammars.

2.1 Logical Theories

We review some key definitions of logical theories used in satisfiability modulo theories (SMT); see [4,11] for more details. A *signature* Σ consists of a set Σ^F of *function* symbols and a set Σ^R of *relation* symbols with an associated *arity*, a non-negative number indicating the number of arguments. We will refer to the 0-arity function symbols as *constants*. Each signature naturally defines a *language* $L(\Sigma)$ of formulas. A first-order model M in Σ, also called Σ-model, is a pair consisting of a set $dom(M)$ called its domain and a mapping $(-)^M$. The mapping assigns to each function symbol $f \in \Sigma^F$ with arity $n \geq 0$ a total function $f^M : dom(M)^n \to dom(M)$, and to each relation $R \in \Sigma^R$ of arity n a set $R^M \subseteq dom(M)^n$.

This mapping defines a natural mapping of formulas $\varphi \in L(\Sigma)$ to truth values $\varphi^M \in \{\text{true}, \text{false}\}$. We write $M \models \varphi$ if $\varphi^M = \text{true}$ and say M *satisfies* φ. For some set Φ of first-order formulas, we say $M \models \Phi$ if $M \models \varphi$ for each $\varphi \in \Phi$. A *theory* $\mathcal{T} \subseteq L(\Sigma)$ is a set of formulas. We say M is a model of $\mathcal{T}$ if $M \models \mathcal{T}$, and use $\text{Mod}(\mathcal{T})$ to denote the set of models of $\mathcal{T}$. A first-order formula φ is *valid* in $\mathcal{T}$ if for all $M \in \text{Mod}(\mathcal{T})$, $M \models \varphi$. A theory is *complete* if for all formulas $\varphi \in L(\Sigma)$ either φ or $\neg\varphi$ is valid.

In this work, we consider the common quantifier-free background theories of SMT solving: propositional logic (SAT), bit-vectors (BV), difference logic (DL), linear real arithmetic (LRA), linear integer (Presburger) arithmetic (LIA), the theory of arrays (AR), and the theory of uninterpreted functions with equality (EUF). For detailed definitions of these theories, see [4,3].

In addition to these "standard" theories, we also consider a family of theories that we term *finite-domain (FD) theories*. Formally, an FD theory is a complete theory that has a finite model. Note that completeness implies that all models of the theory have a finite domain if there is one model with a finite domain (consider the family of sentences that state "there are not n distinct elements", which

can be expressed only using equality, conjunction, and negation). Examples of finite-domain theories are boolean logic or the theory of fixed-length bit-vectors without concatenation. Bit-vector theories with (unrestricted) concatenation allow us to construct arbitrarily many distinct constants and are thus not FD theories. FD theories will be studied more in Section 4.

2.2 Grammars

A *context-free grammar* (CFG) is a tuple $G = (N, T, S, R)$ consisting of a finite set N of nonterminal symbols with a distinguished *start symbol* $S \in N$, a finite set T of terminal symbols, and a finite set R of production rules, which are tuples of the form $(N, (N \cup T)^*)$. Production rules indicate the allowed replacements of non-terminals by sequences over nonterminals and terminals. The *language*, $L(G)$, generated by a context-free grammar is the set of all sequences that contain only terminal symbols that can be derived from the start symbol using the production rules.

Tree grammars are a more restrictive class of grammars. They are defined relative to a *ranked alphabet* $\mathcal{F}$, which is a set of function symbols with an associated arity: $\mathcal{F} \to \mathbb{N}$. We define the set $T_{\mathcal{F}}$ of $\mathcal{F}$*-trees* as the smallest set that includes the symbols $f \in \mathcal{F}$ with arity 0 and further includes for all $f \in \mathcal{F}$ with arity k and elements $t_1, \ldots, t_k \in T_{\mathcal{F}}$ the element $f(t_1, \ldots, t_k)$. A *tree grammar* $G = (N, S, \mathcal{F}, R)$ consists of a set N of non-terminals, a start symbol $S \in N$, a ranked alphabet $\mathcal{F}$, and a set R of production rules. Production rules are of the form $A \to f(t_1, t_2, ..., t_k)$, where $A \in N$, f is in $\mathcal{F}$ and has arity k, and each t_i is in $N \cup T_{\mathcal{F}}$. For a given tree-grammar G we write $L(G)$ for the set of trees produced by G. Any tree grammar can be converted to a CFG by simply treating the right-hand side of any production as a string, rather than a tree. We will thus use $L(G)$ also to denote the set of strings produced by the corresponding CFG.

Let Σ be a signature of a background theory $\mathcal{T}$. We define a tree grammar $G = (N, S, \mathcal{F}, P)$ to be $\mathcal{T}$*-compatible* (or Σ-compatible) if $\mathcal{F} \subseteq \Sigma^F \cup \Sigma^R$ and the arities for all symbols in $\mathcal{F}$ match those in Σ.

A *right-regular grammar*, $G = (N, S, T, P)$, is a CFG in which all productions are of the form $A \to aB$ or $A \to a$, for $A, B \in N$ and $a \in T$. They generate exactly the set of regular languages.

2.3 Syntax-Guided Synthesis

We follow the definition of SyGuS given by Alur et al. [2] and focus on the case that the set of candidate expressions (programs) is generated by a given grammar. Let $\mathcal{T}$ be a background theory, and let $\mathcal{G}$ be a class of grammars. Given a typed function symbol f, a formula φ over the signature of $\mathcal{T}$ along with f, and a grammar $G \in \mathcal{G}$ of expressions in the signature of $\mathcal{T}$ that have the same type as f, the SyGuS problem is to find an expression e generatable by G such that the formula $\varphi[f/e]$ is valid or to determine the absence of such an expression.

Here the notation $\varphi[f/e]$ denotes the formula obtained by substituting e for f in φ.

Even though the problem is given in terms of synthesizing a function f, usually the grammar generates expressions that are *applications* of f.

Example 1. Consider the following example SyGuS problem in linear integer arithmetic. Let the type of the function to synthesize f be $\mathit{int} \times \mathit{int} \mapsto \mathit{int}$ and let the specification be given by the logical formula

$$\varphi_1 : \ f(x,y) = f(y,x) \ \wedge \ f(x,y) \geq x \, .$$

We can restrict the set of expressions $f(x,y)$ to be expressions generated by the grammar below:

$$\begin{aligned} \mathit{Term} &:= x \mid y \mid \mathit{Const} \mid \mathit{ITE}(\mathit{Cond}, \mathit{Term}, \mathit{Term}) \\ \mathit{Cond} &:= \mathit{Term} \leq \mathit{Term} \mid \mathit{Cond} \wedge \mathit{Cond} \mid \neg \mathit{Cond} \mid (\mathit{Cond}) \end{aligned}$$

It is easy to see that a function computing the maximum over x and y, such as $\mathit{ITE}(x \leq y, y, x)$, is a solution to the SyGuS problem. There are, however, other solutions, such as $\mathit{ITE}(7 \leq y \wedge 7 \leq x, \mathit{ITE}(x \leq y, y, x), 10)$. The function computing the sum of x and y would satisfy the specification, but cannot be constructed in the grammar.

3 Equality and Uninterpreted Functions

This section proves that the SyGuS problem for uninterpreted functions with equality (EUF) is undecidable. Our proof is based on a reduction from the Post correspondence problem (PCP) [15]. The reduction relies on a tree grammar that produces only applications of unary function symbols, showing that for EUF, even basic instances of SyGuS are undecidable.

Recall that the Post correspondence problem is to decide, given pairs of finite strings $(s_1, s'_1), (s_2, s'_2), \ldots, (s_n, s'_n)$, the existence of a non-empty sequence of integers $i_1, i_2, \ldots, i_k$ ($1 \leq i_j \leq n$ for all j) such that $s_{i_1} s_{i_2} \ldots s_{i_k} = s'_{i_1} s'_{i_2} \ldots s'_{i_k}$. PCP is undecidable even when the strings only contain two symbols, so we will focus on PCP problems with the alphabet $\{a, b\}$.

As an example consider the Post correspondence problem for the pairs of strings $(s_1, s'_1) = (bb, b)$, $(s_2, s'_2) = (ab, ba)$, and $(s_3, s'_3) = (b, bb)$. The sequence of indices 1 2 2 3 is a solution: $s_1 s_2 s_2 s_3 = bbababb = s'_1 s'_2 s'_2 s'_3$.

Theorem 1. *SyGuS for EUF is undecidable for the class of EUF-compatible tree grammars.*

Proof. Let $P = (s_1, s'_1), (s_2, s'_2), \ldots, (s_n, s'_n)$ be an instance of PCP and let g_a, g_b, g'_a, g'_b, and h be (distinct) unary function symbols. From this PCP instance, we will construct an EUF-compatible tree grammar $G_P = (N, S, T, R)$ and a correctness constraint φ_P.

The terminals of the grammar are "x", "g_a", "g_b", "g'_a", "g'_a", "h", "(", and ")". Given a pair (s_j, s'_j) with $|\sigma| = r$ and $|\sigma'| = r'$, we add the following production rules to R:

$$p_j^S = S \to g_{s_j(1)}(g_{s_j(2)}(\dots(g_{s_j(r)}(g'_{s'_j(1)}(g'_{s'_j(2)}(\dots(g'_{s'_j(r')}(V))\dots))))\dots))$$

$$p_j^V = V \to g_{s_j(1)}(g_{s_j(2)}(\dots(g_{s_j(r)}(g'_{s'_j(1)}(g'_{s'_j(2)}(\dots(g'_{s'_j(r')}(V))\dots))))\dots))\,,$$

where $s_j(i)$ is the i-th symbol in the string s_j. (The production rules occur twice in similar form, once for the start symbol S and once for V, to guarantee that at least one pair of symbols (s_j, s'_j) is produced and so the corresponding PCP solution is non-empty.) To allow the termination of a production sequence we further add the production $p_{end} = V \to x$.

Intuitively, for any sequence of integers $i_1, i_2, \dots, i_k$ such that $s_{i_1} s_{i_2} \dots s_{i_k} = s'_{i_1} s'_{i_2} \dots s'_{i_k}$, we would like φ to be satisfied by the string constructed by following the productions $p^S_{i_1}, p^V_{i_2}, \dots, p^V_{i_k}, p_{end}$. (Since there is at most one non-terminal on the right-hand side of each production rule, this description yields a unique expression.) We now construct the correctness constraint φ such that chains of productions are valid iff the corresponding sequence of indices is a solution to the PCP instance.

Let ψ be the conjunction of the following constraints:

$$\begin{aligned} g_a(g'_b(y)) &= g'_b(g_a(y)) \\ g_b(g'_a(y)) &= g'_a(g_b(y)) \\ h(g_a(g'_a(y))) &= h(y) \\ h(g_b(g'_b(y))) &= h(y) \end{aligned}$$

That is, we require that applications of the functions g_a, g_b commute with applications of the functions g'_a, g'_b (but not necessarily g_a with g_b or g'_a with g'_b). Further, applications of g_a and g'_a and likewise of g_b and g'_b are required to cancel each other when they occur directly inside an application of h. We then define the correctness constraint φ_P to be $\psi \to h(f) = h(x)$. Note that φ is constructed independently of the PCP instance being studied.

The easy direction is to show that every sequence of indices $i_1, i_2, \dots, i_k$ such that $s_{i_1} s_{i_2} \dots s_{i_k} = s'_{i_1} s'_{i_2} \dots s'_{i_k}$ gives us a sequence of productions $p^S_{i_1}, p^V_{i_2}, \dots, p^V_{i_k}$, p_{end} that produces an expression that is valid under every model M. Given any model M that satisfies the antecedent ψ, we can commute the outermost applications of a primed and an unprimed function (of g_a, g_b, g'_a, and g'_b) to the "top". By repeatedly pulling matching applications to the top and eliminating them inside the application of h, we arrive at the term $h(x)$, which satisfies the correctness specification.

We now turn to the other direction. Let $i_1, i_2, \dots, i_k$ be a sequence of indices that is not a solution to the PCP problem. We show that for the expression obtained by the productions $p^S_{i_1}, p^V_{i_2}, \dots, p^V_{i_k}, p_{end}$ there is a model that invalidates the correctness specification $\varphi_P[f/e]$. We construct M such that the antecedent ψ holds for all y, but M does not model $h(e) = h(x)$ for any $e \in L(G_P)$. We define $\mathrm{dom}(M) = \{a, b\}^* \times \{a, b\}^*$ and define g_a, g'_a, g_b, and g'_b to be functions

from $\mathrm{dom}(M)$ to $\mathrm{dom}(M)$. We define these functions as follows, for $\alpha, \beta \in \{a, b\}$ and $s, s' \in \{a, b\}^*$:

$$\begin{aligned} g_\alpha((s, s')) &= (\alpha s, s') \\ g'_\alpha((s, s')) &= (s, \alpha s') \\ h((\alpha s, \beta s')) &= h((s, s')), \text{ if } \alpha = \beta \\ h((\alpha s, \beta s')) &= (\alpha s, \beta s'), \text{ if } \alpha \neq \beta \end{aligned}$$

That is, g_a and g_b add elements to the first list, g'_a and g'_b add elements to the second list, and h removes the common prefix. It is easy to verify that these function definitions satisfy the antecedent ψ.

Assume the variable x in the correctness specifications is instantiated with the tuple $(\varepsilon, \varepsilon)$ of empty strings. If the sequence of letters of a and b produced by the sequence of indices is different in length (say the first is shorter), then we can apply the function definitions above such that the correctness constraint reduces to $(\varepsilon, s) = (\varepsilon, \varepsilon)$ with $s \neq \varepsilon$ and is thus violated. If the sequences of letters a and b are of the same length but have a position in which they are different, we can apply the rules above to eliminate the common prefix and again arrive at a tuple that is different to $(\varepsilon, \varepsilon)$.

We conclude that for any given sequence of indices that is not a solution to P, the corresponding production does not yield an expression that is valid for every model. Since any sequence of productions of the grammar corresponds also to a sequence of indices this concludes the proof. □

3.1 Implications for Arrays

Corollary 1. *SyGuS for AR is undecidable with the class of AR-compatible tree grammars, even when arrays cannot be nested.*

Proof. We can recast the functions in the proof of Theorem 1 as arrays to create a SyGuS problem in AR, which can only be solved if there exists a solution to the original SyGuS problem in EUF. Take any Post-correspondence problem P and grammar G_P as defined above. Let φ be the EUF constraint formula defined above. We can replace any function call $g(a)$ in φ or G_P with the statement $\mathrm{read}(g, a)$ to get a SyGuS problem in AR. It is easy to check that this problem can be solved if, and only if, the original SyGuS problem in EUF is solvable. □

3.2 Extending the Undecidability Result to Regular Grammars

In the proof of Theorem 1, we assumed that the provided grammar was a tree grammar with at most one non-terminal on the right side of any production rule. We may wonder whether an even more restricted grammar could allow the SyGuS problem to become decidable. Consider the right-regular grammars, which generate the regular languages. These grammars cannot even produce the language of well-matched parenthesis, so we cannot expect them to only yield

well-formed strings, that is strings that are expressions in the theory. However, given a right-regular grammar G_{reg} and the grammar of well-formed strings U, we can construct a context-free grammar G' such that $L(G') = L(G_{reg}) \cap L(U)$. Unfortunately, even in this restricted form, we can still specify the grammars used in the above undecidability proofs of EUF and arrays.

Corollary 2. *SyGuS for EUF with right-regular grammars is undecidable.*

Proof. Let $P = (s_1, s'_1), (s_2, s'_2), \ldots, (s_n, s'_n)$ be a PCP instance. We construct G_{reg} with the following productions:

$$\begin{aligned} S \rightarrow &\quad xV \\ V \rightarrow &\quad)V \quad | \quad) \end{aligned}$$

For each (s_i, s'_i) in P, with $|s_i| = r$ and $|s'_j| = r'$, we add the following two production rules to G_{reg}:

$$S \rightarrow \quad g_{s_j(1)}(g_{s_j(2)}(\ldots(g_{s_j(r)}(g'_{s'_j(1)}(g'_{s'_j(2)}(\ldots(g'_{s'_j(r')}(\ V$$

$$V \rightarrow \quad g_{s_j(1)}(g_{s_j(2)}(\ldots(g_{s_j(r)}(g'_{s'_j(1)}(g'_{s'_j(2)}(\ldots(g'_{s'_j(r')}(\ V$$

We can see that the set of well-formed expressions from G_{reg} precisely matches those of grammar G_P as defined in Theorem 1. □

Corollary 3. *SyGuS for AR is undecidable with regular languages, even when arrays cannot be nested.*

Proof. Let $P = (s_1, s'_1), (s_2, s'_2), \ldots, (s_n, s'_n)$ be a PCP instance. We create a grammar G_{reg} and add the following productions:

$$\begin{aligned} S \rightarrow &\quad xV \\ V \rightarrow &\quad)V \quad | \quad) \end{aligned}$$

For each (s_i, s'_i) in P, with $|s_i| = r$ and $|s'_j| = r'$, we add the following transitions to G_{reg}:

$$\begin{aligned} S \rightarrow &\quad read(g_{s_j(1)}, read(g_{s_j(2)}, \ldots, read(g_{s_j(r)}, \\ &\qquad read(g'_{s'_j(1)}, read(g'_{s'_j(2)}, \ldots, read(g'_{s'_j(r')},\ V \\ V \rightarrow &\quad read(g_{s_j(1)}, read(g_{s_j(2)}, \ldots, read(g_{s_j(r)}, \\ &\qquad read(g'_{s'_j(1)}, read(g'_{s'_j(2)}, \ldots, read(g'_{s'_j(r')},\ V \end{aligned}$$

The well-formed expressions of G_{reg} precisely match the expressions from Corollary 1. □

3.3 On Well-Formed and Ill-Formed Expressions

When we allow the grammar in a SyGuS problem to produce ill-formed expressions and consider the intersection with the expressions of the background theory, as in the previous subsection, we can derive a very strong proof of undecidability. The proof only relies on the well-formedness of expressions, and holds independently of the interpretation of the function symbols. That is, the proof also holds for theories other than EUF.

Theorem 2. *SyGuS for CFGs is undecidable for any theory $\mathcal{T}$ with domain D containing two function symbols $f : D^m \to D$ and $g : D^n \to D$ of different non-zero arities ($m \neq n$ and $m, n > 0$).*

Proof. This proof is by reduction from the Post-correspondence problem. Let h and g be function symbols in $sig(\mathcal{T})$ of arity m and n, respectively, where $m \neq n$ and $m, n > 0$. Let $P = (s_1, s'_1), (s_2, s'_2), \ldots, (s_k, s'_k)$ be a PCP instance with alphabet $\{a, b\}$. We construct a grammar $G_P = (N, S, T, R)$ so that sequences of indices that are a solution to the PCP problem correspond to derivations in the grammar that produce a well-formed expression, and where sequences of indices that are not a solution to the PCP problem correspond to derivations in the grammar that produce an ill-formed expression.

For each pair (s_i, s'_i) in P we construct two terminals σ_i and σ'_i in T. For each position j in the sequence $s_i \in \{a, b\}^*$ we define $\sigma_i(j) =$ "$f($" if $s_i(j) = a$ and $\sigma_i(j) =$ "$g($" if $s_i(j) = b$. Correspondingly, we define for each position j in $s'_i \in \{a, b\}^*$ that

$$\sigma'_i(r + 1 - j) = (\text{“}, x\text{”})^{m-1} + \text{“)”} \quad \text{if } s'_i(j) = a \quad \text{and}$$

$$\sigma'_i(r + 1 - j) = (\text{“}, x\text{”})^{n-1} + \text{“)”} \quad \text{if } s'_i(j) = b\,,$$

where $r = |s'_i|$ is the length of the word s'_i and x is some variable. That is, the terminals σ'_i are lists of lists of arguments of length $m - 1$ and $n - 1$ that are supposed to match the function calls of f and g, respectively. We are only providing $m - 1$ and $n - 1$ arguments, respectively, as the first argument will (except for the innermost level) be the next function call. Within the terminals σ'_i the argument lists occur in *reverse order* compared to the strings s'_i.

We add two nonterminals S and V and for each pair of strings (s_i, s'_i) in P we add one production rule for both nonterminals:

$$S \to \sigma_i V \sigma'_i \quad \text{and} \quad V \to \sigma_i V \sigma'_i$$

That is, the production rules produce a string over opening parts of function calls (σ_i) matching, in reverse order, a string of argument lists (σ'_i) which each are reversed internally. Finally, we add one production rule $V \to$ "x", which provides the "missing" argument for the innermost function call. (The need for two nonterminals comes from the fact that we have to disallow the derivation yielding only x.)

We can then define the correctness constraint φ as $f = f$, where f must be replaced by some expression in G_P. For this correctness constraint, each syntactically correct replacement of f is valid in every theory. Thus, this SyGuS problem has a solution, if and only if the PCP instance P has a solution. □

Example 2. Let $(s_1, s_1') = (bb, b)$, and $(s_2, s_2') = (a, ba)$. Then $s_1 s_2 = bba = s_1' s_2'$ is a solution to the Post-correspondence problem. Let f and g be functions in the background theory $\mathcal{T}$ with arities 1 and 2, respectively. We can associate f with a and g with b. The grammar G_P can produce the expression $g(g(f(x), x), x)$, which corresponds to the solution *bba*.

4 Finite-Domain Theories

Finite-domain theories include, for example, theories of bit-vectors without the concatenation operation (or a bounded version of concatenation). In this section we give a generic algorithm for any complete finite-domain theory for which validity is decidable. Let $\mathcal{T}$ be a such a theory and let M be a model of $\mathcal{T}$ with a finite domain $dom(M)$. Assume without loss of generality that for every element $c \in dom(M)$ there is a constant f in M such that $f^M = c$.

We consider a SyGuS problem with a correctness specification φ in theory $\mathcal{T}$, a function symbol f to synthesize, and a tree grammar $G = (N, S, \mathcal{F}, P)$ generating the set of candidate expressions. Let $\mathbf{x} = x_0, x_1, \ldots, x_r$ be the variables (i.e. the uninterpreted symbols) occurring in φ. The expression e generated by G to replace f can be seen as a function mapping $\mathbf{x}$ to an element in $dom(M)$. Even though there are only finitely many candidate functions if the domain of M is finite, it can be non-trivial to determine which functions can be generated by G. In the following, we describe an algorithm that iteratively determines the set of functions that can be generated by each non-terminal in the grammar G.

For each $V \in N$, we maintain a set E_V of expressions e. In each iteration and for each production rule $V \to f(t_1, \ldots, t_k)$ for V in G, we consider the expressions $f(t_1^*, \ldots, t_k^*)$ where $t_i^* = t_i$ if t_i is an expression (i.e. $t_i \in T_{\mathcal{F}}$) and $t_i^* \in E_{V'}$ if t_i is a non-terminal V'. Given such an expression e, we compute the function table, that is the result of $e[\mathbf{x}/\mathbf{c}]$ for each $\mathbf{c} \in dom(M)^r$, compare it to the function table of the expressions currently in E_V. Decidability of validity for $\mathcal{T}$ guarantees that this operation is decidable. If e represents a new function, we add it to the set E_V.

The algorithm terminates, if for no nonterminal V the set E_V changed for one iteration. As there are only finitely many functions from $dom(M)^r$ to $dom(M)$ and the sets E_V grow monotonously, the algorithm eventually terminates. To determine the answer to the SyGuS problem, we then check whether there is an expression e in E_S, for which $\varphi[f/e]$ is valid.

Theorem 3. *Let $\mathcal{T}$ be a complete theory for which validity is decidable and which has a finite-domain model M. The SyGuS problem for $\mathcal{T}$ and $\mathcal{T}$-compatible tree grammars is decidable.*

Iteration#	E_S	E_A	E_B
1	none	x	y
2	$x \oplus y$	$\neg y$	none
3	$\neg y \oplus y \equiv \top$	none	$(x \oplus y) \oplus \neg y \equiv \neg x$ ~~$(x \oplus y) \oplus x \equiv y$~~
4	$\neg y \oplus \neg x$ ~~$\neg x \oplus x \equiv \top$~~	~~$\neg\neg x \equiv x$~~ none	~~$\top \oplus \neg y \equiv y$~~ ~~$\top \oplus x \equiv \neg x$~~ none
5	none	none	~~$(\neg y \oplus \neg x) \oplus \neg y \equiv \neg x$~~ ~~$(\neg y \oplus \neg x) \oplus x \equiv y$~~ none

Table 2. This table shows the expressions added to the sets E_S, E_A, and E_B when we apply the algorithm to the SyGuS problem in Example 3. For readability, we simplify the expressions, indicated by the symbol '$\equiv$'. Expressions that are syntactically new, but do not represent a new function are struck out. When no new function is added, "none" is written in the cell.

Example 3. Consider the SyGuS problem over boolean expressions with the correctness specification $\varphi = x \oplus f$, where $\oplus$ denotes the XOR operation and f is the function symbol to synthesize from following tree grammar (we use infix operators for readability):

$$S \to (A \oplus B)$$
$$A \to \neg B \mid x$$
$$B \to (S \oplus A) \mid y$$

The grammar generates boolean functions of variables x and y and the updates to E_A, E_B, and E_S during each iteration of the proposed algorithm are given in Table 2. The next step in the algorithm is to determine if any of the three expressions $E_S = \{x \oplus y, \neg y \oplus y, \neg y \oplus \neg x\}$ make the formula $\varphi[f/e]$ valid, which is not the case.

5 Bit-Vectors

In this section, we show that the SyGuS problem for the theory of bit-vectors is undecidable - even when we restrict the problem to tree grammars. The proof makes use of the fact that we can construct (bit-)strings with the concatenation operation and can compare arbitrarily large strings with the equality operation. This enables us to encode problem of determining if the languages of CFGs with no ε-transitions have non-empty intersection, which is undecidable [12].

Theorem 4. *The SyGuS problem for the theory of bit-vectors is undecidable for both the class of context-free grammars and the class of BV-compatible tree grammars.*

Proof. We start with the proof for the class of context-free grammars. Given two context-free grammars $G_1 = (N_1, S_1, T_1, R_1)$ and $G_2 = (N_1, S_2, T_2, R_2)$, we define a SyGuS problem with a single context-free grammar $G = (N, S, T, R)$ that has a solution iff the intersection of G_1 and G_2 is not empty. The proof idea is to express the intersection of the two grammars as the equality between two expressions, each generated by one of the grammars. The new grammar thus starts with the following production rule:

$$S \to S_1 = S_2$$

We then have to translate the grammars G_1 and G_2 into grammars G_1' and G_2' that produce expressions in the bit vector theory instead of arbitrary strings over their alphabets. There is a string produced by both G_1 and G_2 if and only if the constructed grammars G_1' and G_2' can produce a pair of equal expressions. We achieve this by encoding each letter as a bit string of the fixed length $1+\log_2 |T_1 \cup T_2|$, and by intercalating concatenation operators (@) in the production rules: We encode each production rule (N, P) with $P = p_1 p_2 \dots p_n$ as (N, P') with $P' = p_1' @ p_2' @ \dots @ p_n'$, where $p_i' = p_i$ if $p_i \in N$, and otherwise p_i' are the fixed-length encodings of the terminal symbols. We then define $N = S \dot\cup N_1' \dot\cup N_2'$, $T = \{0, 1, @, =\}$, and $R = R_1' \cup R_2' \cup \{(S, S_1' = S_2')\}$. The correctness constraint φ of our SyGus problem then only states $\varphi \coloneqq \neg f$, where $f : \mathbb{B}$ is the function symbol, a constant, to synthesize. As each character in the alphabets of the context-free grammars was encoded using bit vectors the same length, the comparison of the bit vectors is equivalent to the comparison between the strings of characters of the grammars G_1 and G_2 and the SyGuS problem has a solution if and only if the intersection of the languages of the context-free grammars G_1 and G_2 is empty.

Note that the context-free grammar G can also be interpreted as a BV-compatible tree grammar, where BV is the theory of bit-vectors. Although it is efficiently decidable whether two tree grammars produce a common tree, the expressions produced by the tree-interpretation of G_1 and G_2 will be equivalent as long as their leaves are equivalent. Thus, the equality of the expression trees in the interpretation of the bit vector theory still coincides with the intersection of the given context-free grammars G_1 and G_2. □

We only used the concatenation operation of the bit-vector theory for the proof. That is, SyGuS is even undecidable for fragments of the theory of bit-vectors for which basic decision problems are easier than the general class; for example, the theory of *fixed-sized bit-vectors with extraction and composition* [6] for which satisfiability of conjunctions of atomic constraints is polynomial-time solvable unlike the general case which is NP-hard.

Remark 1. This proof only relies on the comparison of arbitrarily large values in the underlying logical theory. It may thus be possible to extend the proof to other theories involving numbers, such as LIA, LRA, and difference logic. The problem here is that these proofs tend to depend on syntactical sugar. Consider the case of LIA. If the signature allows us to use arbitrary integer constants, such

as 42, it is simple to translate the proof above into a proof of undecidability of SyGuS for LIA and CFGs. For the standard signature of LIA, however, which just includes the integer constants 0 and 1 (larger integers can then be expressed as the repeated addition of the constant 1) the proof scheme above does not apply.

6 Other Background Theories

In this section, we remark on the decidability for some other classes of SyGuS problems. These results are straightforward, but the classes do occur in practice, and so they are worth mentioning.

Linear real arithmetic (LRA) with arbitrary affine expressions. Consider the family of SyGuS problems where:

i) the specification φ is a Boolean combination of linear constraints over real-valued variables $\boldsymbol{x} = x_1, x_2, \ldots, x_n$ and applications of the function f to be synthesized. For simplicity, we assume a single function f of arity n; the arguments below generalize.
ii) The grammar G is the one generating *arbitrary affine expressions* over $\boldsymbol{x}$ to replace for applications of f. Thus, the application $f(\boldsymbol{t})$, where $\boldsymbol{t} = t_1, t_2, \ldots, t_n$ is a vector of LRA terms, is replaced by an expression of the form $a_0 + \sum_{i=1}^{n} a_i t_i$.

Thus, for a fixed set of variables $\boldsymbol{x}$ there is a fixed grammar for all formulas φ.

This case commonly arises in invariant synthesis when the invariant is hypothesized to be an affine constraint over terms in a program. In this case, the solution of the SyGuS problem reduces to solving the $\exists\forall$ SMT problem

$$\exists a_0, a_1, \ldots, a_n \,.\, \forall x_1 x_2 \ldots x_n \,.\, \left(\varphi[f(\boldsymbol{t})/a_0 + \sum_{i=1}^{n} a_i t_i]\right)$$

which reduces to a formula with first-order quantification over real variables. Since the theory of linear real arithmetic admits quantifier elimination, the problem is solvable using any of a number of quantifier elimination techniques, including classic methods such as Fourier-Motzkin elimination [7] and the method of Ferrante and Rackhoff [9], as well as more recent methods for solving exists-forall SMT problems (e.g., [8]).

This decidability result continues to hold for grammars that generate bounded-depth conditional affine expressions. However, the case of unbounded-depth conditional affine expressions is still open, to our knowledge.

A similar reduction, for the case of affine expressions, can be performed for linear arithmetic over the integers (LIA), requiring quantifier elimination for Presburger arithmetic. Thus, this case is also decidable.

Finite-precision bit-vector arithmetic (BV) with arbitrary bit-vector functions. Consider the family of SyGuS problems where:

i) the specification φ is an arbitrary formula in the quantifier-free theory of finite-precision bit-vector arithmetic [4,3] over a collection of k bit-vector variables whose cumulative bit-width is w. Let f be a bit-vector function to be synthesized with output bit-width m.
ii) The grammar G is the one generating *arbitrary bit-vector expressions* over these k variables, using all the operators defined in the theory. In other words, G imposes no major syntactic restriction on the form of the bit-vector function f.

Thus, for a fixed set of bit-vector variables there is a fixed grammar for all formulas φ.

This class of SyGuS problems has been studied as the synthesis of "bitvector programs" (in applications such as code optimization and program deobfuscation) from components (bit-vector operators and constants) [13,10]. It is easy to see that this class is decidable. A simplistic (but not very efficient) way to solve it is to enumerate all 2^{m2^w} possible semantically-distinct bitvector functions over the k variables and check, via a SMT query, whether each, when substituted for f will make the resulting formula valid.

7 Discussion

In this paper, we have presented a first theoretical analysis of the SyGuS problem, focusing on its decidability for various combinations of logical theories and grammars. The main results of the paper are summarized in Table 1, augmented by the decidability of the simple but common SyGuS classes described in Section 6. We conclude with a few remarks about the results, connections between them, and their relevance in practice.

Consider the theory of finite-precision bit-vector arithmetic (BV). We have seen in Section 6 that the SyGuS problem is decidable when the logical formula is an arbitrary BV formula and the grammar allows the function to be replaced by any bit-vector function over the constants in the formula. However, we have also seen that the SyGuS problem is undecidable when we restrict the set of functions to those that can be generated by an arbitrary context-free grammar (see Section 5). These results may seem to contradict the intuition (stated in Sec. 1) that syntax guidance restricts the search space for synthesis and thus makes the problem easier to solve. We thus have to be careful which classes of grammars we pick to restrict SyGuS problems.

For future work, it would be good to study the LIA and LRA background theories in more detail. In particular, we would like to determine if these theories are decidable when grammars are provided, and whether the use of conditionals without bounding expression tree depth affects the decidability. Further, for SyGuS classes that are decidable, it would be useful to perform a more fine-grained characterization of problem complexity, especially with regard to special classes of grammars.

Acknowledgements. This work was supported in part by NSF Expeditions project CCF-1139138.

References

1. R. Alur, M. Martin, M. Raghothaman, C. Stergiou, S. Tripakis, and A. Udupa. Synthesizing Finite-state Protocols from Scenarios and Requirements. In *Haifa Verification Conference*, volume 8855 of *LNCS*. Springer, 2014.
2. Rajeev Alur, Rastislav Bodik, Garvit Juniwal, Milo M. K. Martin, Mukund Raghothaman, Sanjit A. Seshia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak, and Abhishek Udupa. Syntax-guided synthesis. In *Proceedings of the IEEE International Conference on Formal Methods in Computer-Aided Design (FMCAD)*, pages 1–17, October 2013.
3. Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The SMT-LIB standard: Version 2.5, 2010.
4. Clark Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli. Satisfiability modulo theories. In Armin Biere, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, volume 4, chapter 8. IOS Press, 2009.
5. Michael Colón, Sriram Sankaranarayanan, and Henny Sipma. Linear invariant generation using non-linear constraint solving. In *Computer Aided Verification, 15th International Conference*, pages 420–432, 2003.
6. David Cyrluk, Oliver Möller, and Harald Rueß. An efficient decision procedure for the theory of fixed-sized bit-vectors. In *Computer Aided Verification*, pages 60–71. Springer, 1997.
7. G. B. Dantzig and B. C. Eaves. Fourier-Motzkin elimination and its dual. *Journal of Combinatorial Theory A*, 14:288–297, 1973.
8. Bruno Dutertre. Solving exists/forall problems with Yices. In *In International Workshop on Satisfiability Modulo Theories (SMT)*, 2015.
9. Jeanne Ferrante and Charles Rackoff. A decision procedure for the first order theory of real addition with order. *SIAM Journal of Computing*, 4(1):69–76, 1975.
10. Sumit Gulwani, Susmit Jha, Ashish Tiwari, and Ramarathnam Venkatesan. Synthesis of loop-free programs. *SIGPLAN Notices*, 46:62–73, June 2011.
11. Wilfrid Hodges. *A shorter model theory*. Cambridge university press, 1997.
12. John E Hopcroft and Jeffrey D Ullman. *Introduction to automata theory, languages, and computation*. Pearson Education India, 1979.
13. Susmit Jha, Sumit Gulwani, Sanjit A. Seshia, and Ashish Tiwari. Oracle-guided component-based program synthesis. In *ICSE*, pages 215–224, 2010.
14. Z. Manna and R. Waldinger. A deductive approach to program synthesis. *ACM TOPLAS*, 2(1):90–121, 1980.
15. Emil L Post. A variant of a recursively unsolvable problem. *Bulletin of the American Mathematical Society*, 52(4):264–268, 1946.
16. Armando Solar-Lezama, Rodric Rabbah, Rastislav Bodík, and Kemal Ebcioglu. Programming by sketching for bit-streaming programs. In *PLDI*, pages 281–294, 2005.
17. Armando Solar-Lezama, Liviu Tancau, Rastislav Bodík, Sanjit A. Seshia, and Vijay Saraswat. Combinatorial sketching for finite programs. In *ASPLOS*, pages 404–415, 2006.
18. Abhishek Udupa, Arun Raghavan, Jyotirmoy V. Deshmukh, Sela Mador-Haim, Milo M.K. Martin, and Rajeev Alur. TRANSIT: Specifying protocols with concolic snippets. In *Proceedings of the* 34^{th} *ACM SIGPLAN conference on Programming Language Design and Implementation*, pages 287–296, 2013.